
BIG

Revolutionizing how customers build, deploy and operate highly available and scalable services

Application Platform Infrastructure Team
Windows Server Product Group
Microsoft Corporation

June 18, 2002

Agenda

- Vision
- Architecture
 - *Hardware Reference Platform*
 - *Resource Management*
 - *iBIG*
 - *Service Definition Model (SDM)*
 - *Operations Logic*
- Customers and Scenarios
- Wrap-up

• Enable

- **development** of distributed, scalable and highly available services using Visual Studio and **reusable** building blocks like SQL, IIS.
- **deployment** across a set of **abstracted** hardware resources that are automatically allocated, purposed and configured
- **lower cost of ownership** through automation by codifying operational best practices to control service availability and growth
- **procurement** of standardized data center hardware that leverages **commodity economics**

BIG is a revolutionary re-think of the way highly available and scalable services are built, deployed, operated

BIG Will Deliver a Services Platform

- **Hardware reference platform that aggregates commodity hardware to build a single large computer that we call the BIG Computer**
 - *Includes interconnected servers, network devices, and storage*
- **Hardware abstraction layer that virtualizes resources**
 - *Enables dynamic hardware binding and re-deployment and automated network configuration*
- **Service Definition Model (SDM) for developers to describe an entire service**
 - *Enables developers to rapidly build new services using highly available SQL, IIS and other reusable building block components*
 - *Requires incremental changes to BIG enable existing services*
- **Highly available runtime that supports the SDM**
 - *Enables hosting multiple scalable services inside the BIG Computer*
- **Operations logic framework for automating operational best practices**
 - *Enables policy expression and enforcement*

Enterprise

BIG Ecosystem

Enterprise



OEM



BIG Ecosystem

Enterprise

OEM



BIG Ecosystem

Developer



Enterprise



OEM



BIG Ecosystem

Developer



ASP.NET

Enterprise



OEM



BIG Ecosystem

Developer



ASP.NET



SQL

Enterprise



OEM



Community


Enterprise

EM



BIG Services Platform

Big Computer



ASP.NET

SQL

Operations Log

BIG Ecosystem

Developer

Enterprise


EM



eBusiness

BIG Services Platform

Big Computer



ASP.NET

Operations Log

BIG Ecosystem

Developer

Enterprise

OEM



BIG Services Platform

BIG Services Platform



ACP IT

IT

Operations Log

BIG Ecosystem

Developer


Enterprise

CEM



BIG Services Platform

BI a Computer



ACP BT

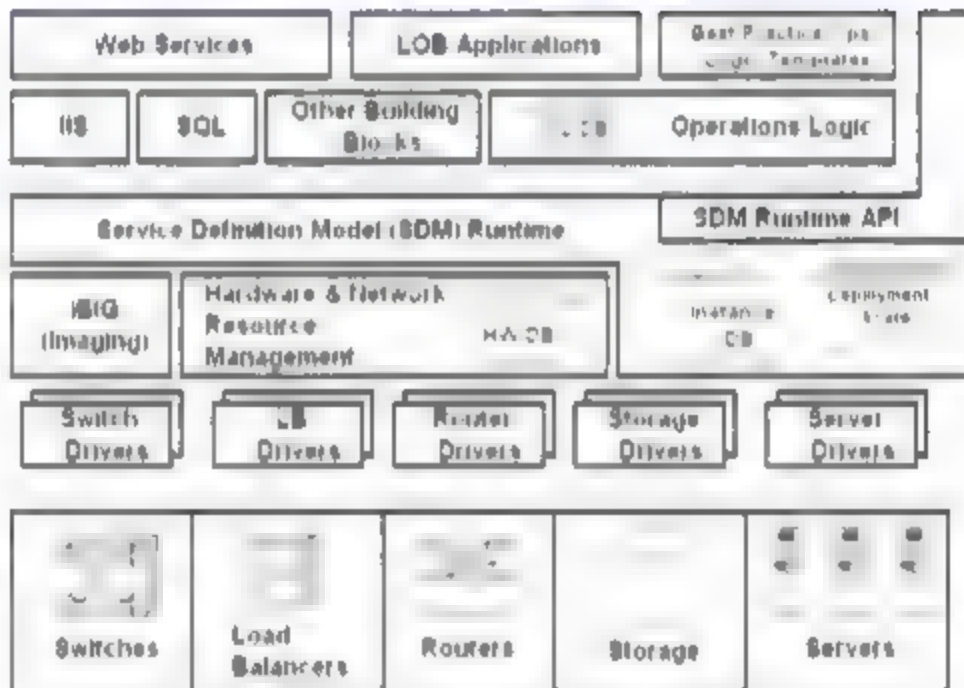
Operations Log

BIG Services Platform Architecture

Services
Components
& Operations

BIG Services
Platform


BIG Computer
(Hardware
Reference)



Hardware Reference Platform Initiative

BIG Computer

Hardware
Reference)




Hardware Reference Platform (BIG Computer)

Motivation

- Reduces the cost of design, test and operations
- Limits number of hardware devices to support
- Constraints the network topology and enables automation of network configuration
- Customers adhere to BIG taking over the entire data center P/E (DHCP, DNS, network switches)



IP Gateway



Internal Network



Hardware Building Blocks



Examples of Current Products that can be Inside a BIG Computer

BL type



Compaq ProLiant BL e-series



enclosure 16 blades



enclosure 14 blades

IBM BladeCenter

Dell PowerEdge 1855MC



enclosure 16 blades

HP bc 1100



enclosure 16 blades




Patric blades

-
-
-

Infrastructure Services

Big Services
Platform



Examples of Current Products that can be Inside a BIG Computer

BL 1100



Compaq ProLiant BL 1100s



enclosure 1 blades



enclosure 14 blades

IBM BladeCenter

Dell PowerEdge 1855MC



enclosure 6 blades

HP bc1150



enclosure 10 blades



Fujitsu blades



BIG Hardware Resource Discovery and Management

- Dynamic or static hardware discovery
- Resource drivers are bound to hardware devices and express logical resources for allocation
- Allocation translates logical resource request to physical resource availability

Right-click on the

Power icon in the



Properties

- Power
- Network
- Storage
- Processor
- Memory
- Location


Physical Resource Graph

Physical Resource Database

Network Management within the BIG Computer

- BIG Computer defines an abstraction layer for network management
- BIG programs the network switch hardware based on application through software
- VLANs provide isolation
 - Management VLAN - for management
 - External VLAN - for external access
 - MyService1 VLAN - for MyService1
 - MyService2 VLAN - for MyService2
 - System VLAN - for system
- Network components
 - Load-Balancer
 - Firewall
 - NAT
 - BIG Runtime
 - Resource Mgr
 - I/O

BIG
Computer



IBIG Features

- **Base Deployment Services**

- Remote deployment of applications and updates
- Remote deployment of operating system updates
- Remote deployment of device drivers

- **Image Deployment and Management**

- Remote deployment of operating system images
- Remote deployment of application images

- **Multiple Device Management (MDM)**

- Remote deployment of device management policies
- Remote deployment of device management updates
- Remote deployment of device management settings

- **Supports Windows 2000 and .NET Server targets**

IBIG Architecture

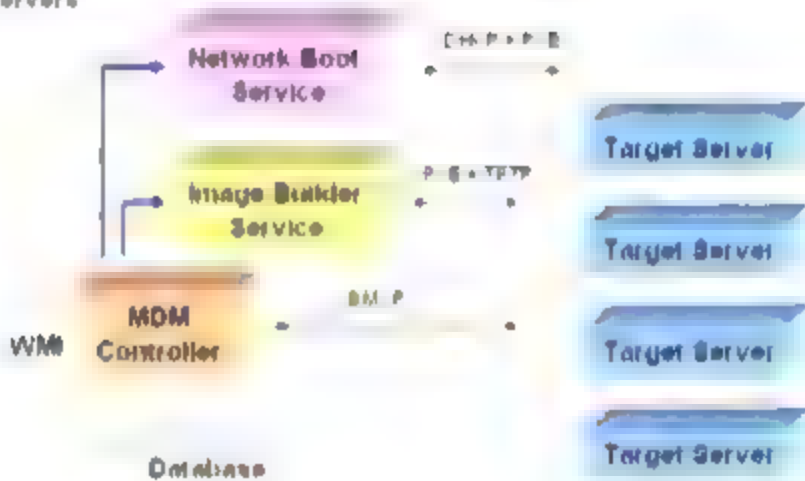
Services can be relocated
or on separate servers

Target servers run IBIG
agent to communicate
with controller

Agent on:

mmio
line

MMC UI



Controller runs only on .NET Enterprise Server
WMI is the interface to the controller

- Windows 2000
- .NET Server

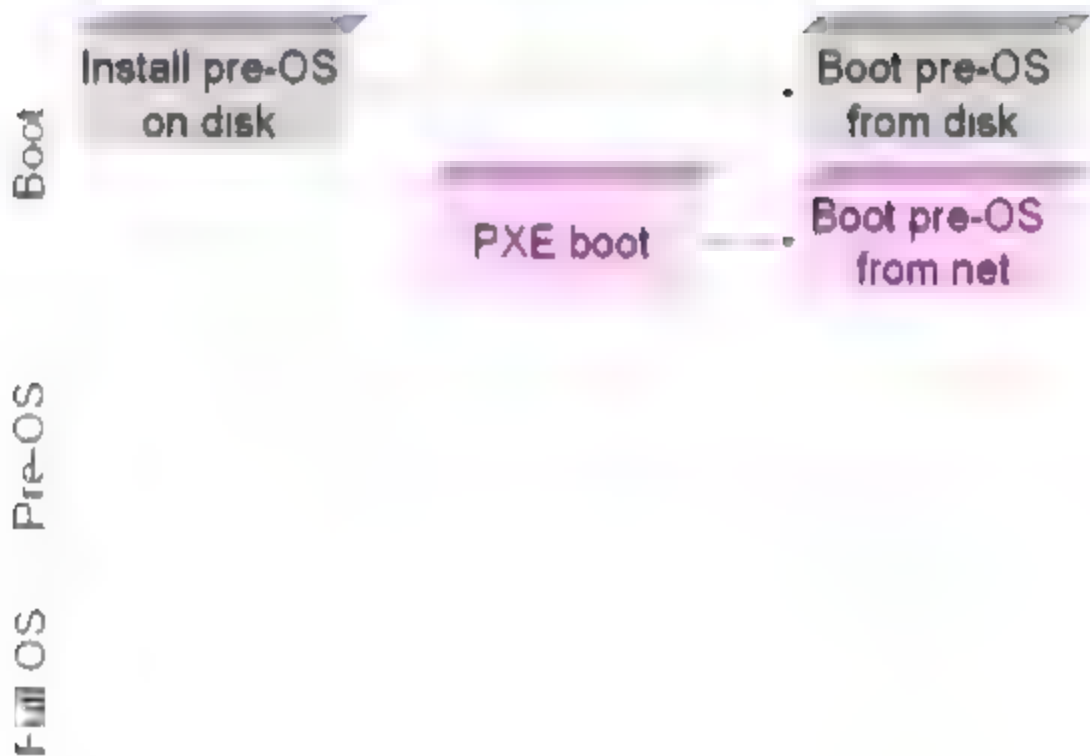
Server Purposing

Boot

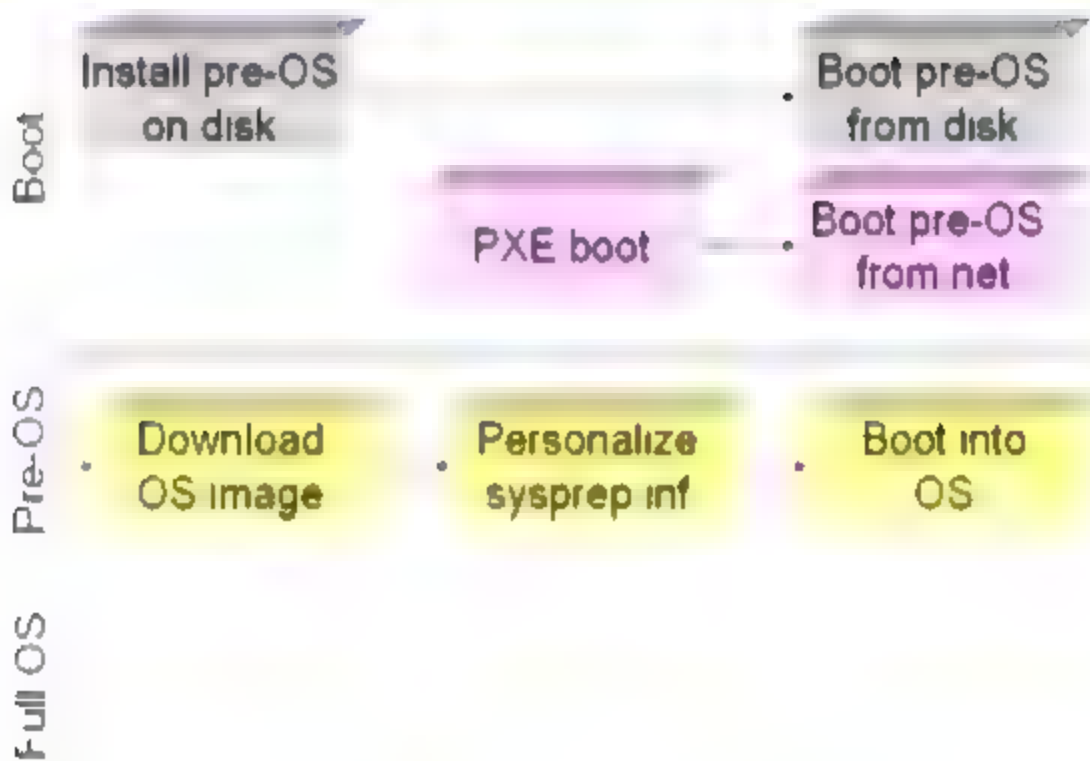
Pre-OS

Full OS

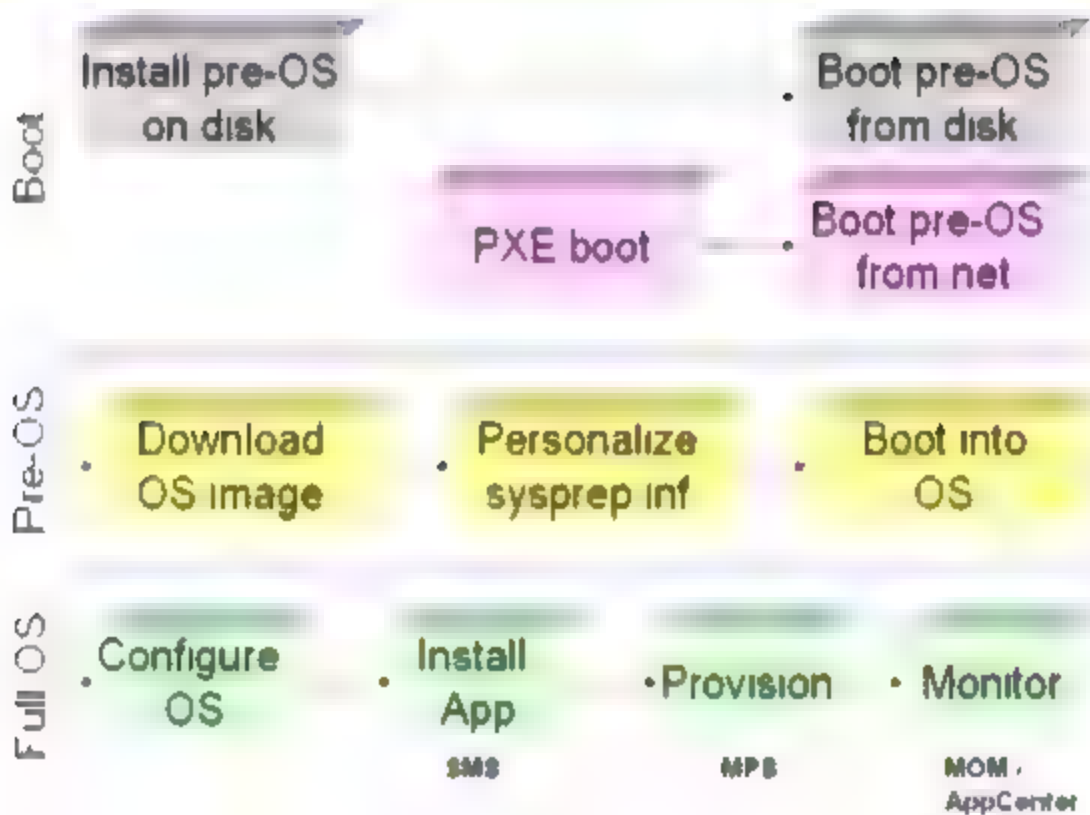
Server Purposing



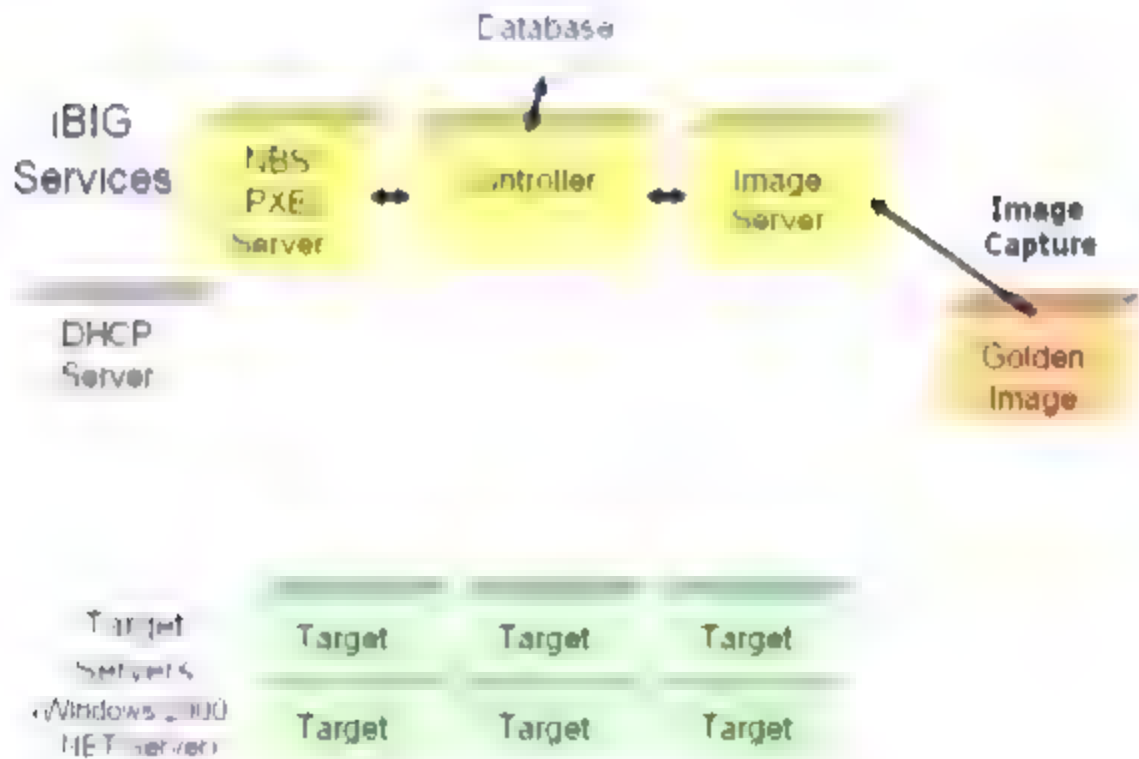
Server Purposing



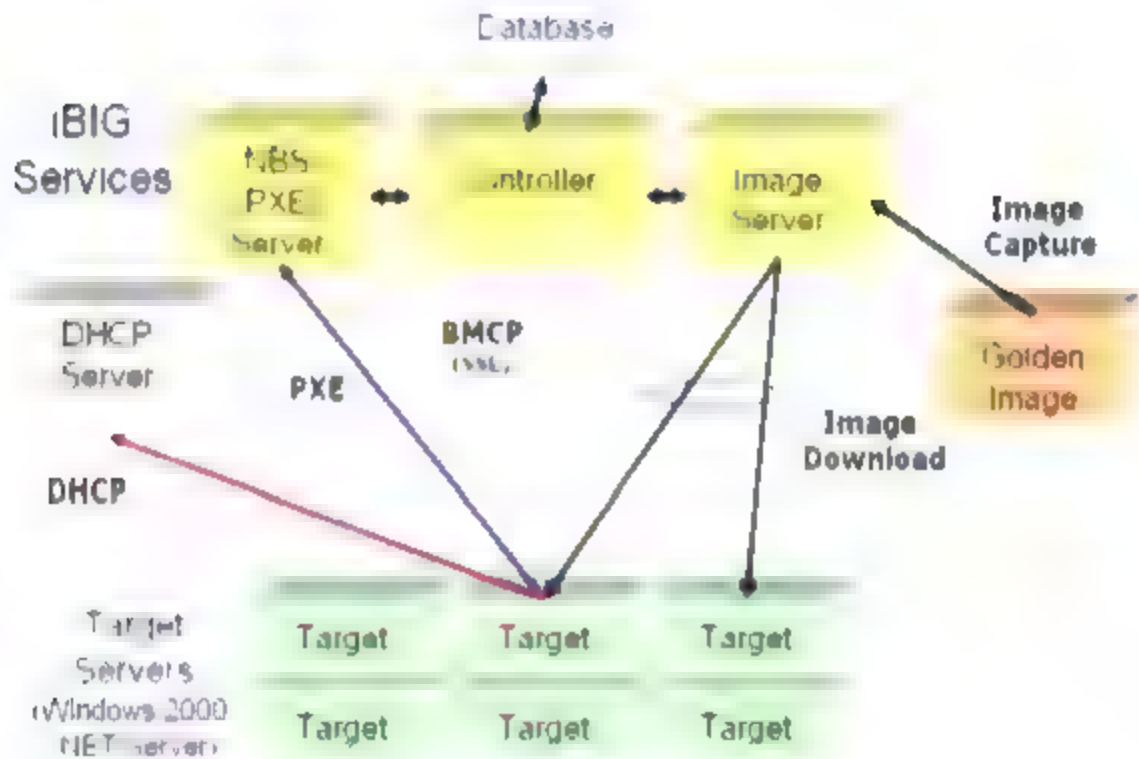
Server Purposing



Remote Boot and Imaging



Remote Boot and Imaging



Service Definition Model (SDM)

- The programmatic description of the entire service

- Declarative \Rightarrow "What" not "How"

- \Rightarrow "What" \Rightarrow service structure, "How" \Rightarrow implementation

- Platform/framework \Rightarrow "How" not "What"

- Component-based \Rightarrow "How" not "What"

- SDML is the declarative language for defining Service Definition Models

Components, Ports and Wires

- **Components** are units of implementation, deployment and operations
 - `Component` is an abstract class
 - `Component` is a `NamedElement`
 - `Component` is a `Deployable`
- **Ports** are names (service access points) with an associated type (protocol)
 - `Port` is an abstract class
 - `Port` is a `NamedElement`
- **Wires** are the permissible bindings between ports
 - `Wire` is an abstract class

Components, Ports and Wires

- **Components** are units of implementation, deployment and operations
 - `Component` is an abstract class
 - `Component` is a `NamedElement`
 - `Component` is a `Deployable`
- **Ports** are names (service access points) with an associated type (protocol)
 - `Port` is an abstract class
 - `Port` is a `NamedElement`
- **Wires** are the permissible bindings between ports
 - `Wire` is an abstract class


SDML Example: MyService.sdml

```
using System;  
using System SQL;  
using System IIS;  
assembly name MyService
```

```
componenttype MyService  
{  
  componenttype MyBackend {  
    SQLDatabase {  
      implementation MySQL MyCLRApp  
    }  
  }  
}
```

```
componenttype MyService  
{  
  componenttype MyBackend {  
    port http = fo http;  
    wire TDS tds {  
      fo catalog  
      fo sql  
    }  
    implementation MyService = MyCLRApp  
  }  
}
```

MyService



Components, Ports and Wires

- **Components** are units of implementation, deployment and operations
 - `Component` is an abstract class
 - `Component` is a `NamedElement`
 - `Component` is a `Deployable`
- **Ports** are names (service access points) with an associated type (protocol)
 - `Port` is an abstract class
 - `Port` is a `NamedElement`
- **Wires** are the permissible bindings between ports
 - `Wire` is an abstract class

BIG Hardware Resource Discovery and Management

- Dynamic or static hardware discovery
- Resources (drivers) are bound to hardware (devices) and expose logical resources for allocation
- Allocation translates logical resource request to physical resource availability

Room 424614

Power 300480



Properties

- Power
- Network
- Storage
- Processor
- Memory
- Location

Physical Resource Graph

Physical Resource Database

Components, Ports and Wires


- Components are units of implementation, deployment and operations
 - `Component` is an abstract class
 - `Component` is a `NamedElement`
 - `Component` is a `Deployable`
- Ports are names (service access points) with an associated type (protocol)
 - `Port` is an abstract class
 - `Port` is a `NamedElement`
- Wires are the permissible bindings between ports
 - `Wire` is an abstract class

SDML Example: MyService.sdml

```
using System;  
using System.SQL;  
using System.IIS;  
assembly name MyService
```

```
componenttype MyFrontEnd  
  ASPApplication {  
    SQLConnection {  
      ...  
    }  
  }  
  
componenttype MyBackEnd  
  SQLDatabase {  
    implementation MySQL.MyCLRApp  
  }  
}
```

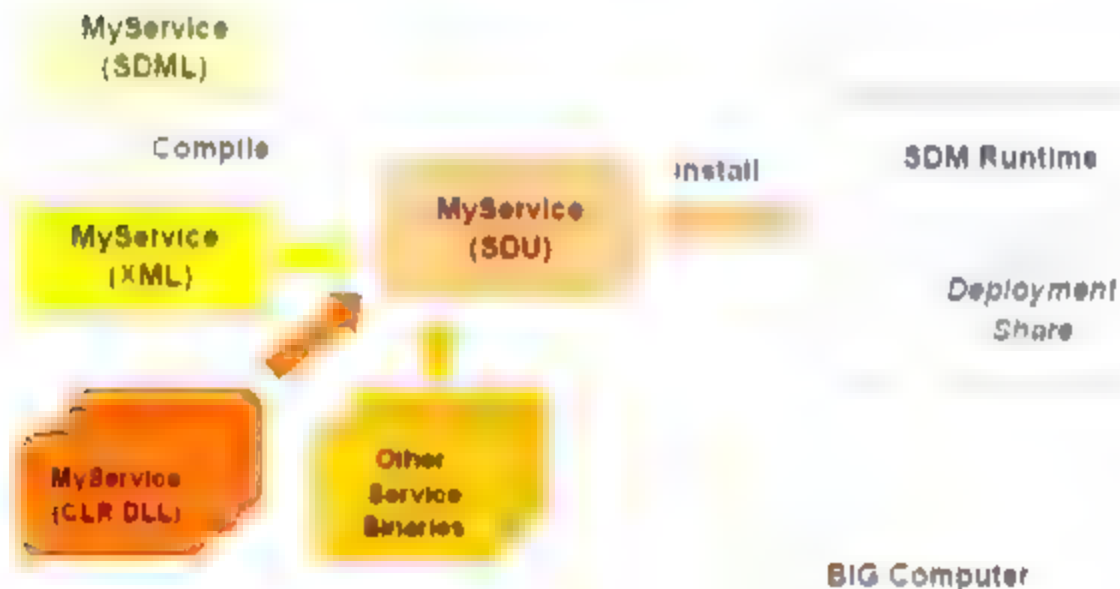
```
componenttype MyService  
{  
  * componenttype MyFrontEnd {  
  * componenttype MyBackEnd {  
    port http * to http,  
    wire TDS {  
      to catalog  
      to sql  
    }  
  }  
  implementation MyService = MyCLRApp  
}
```



Service Deployment Unit (SDU)

- Encapsulates all the pieces that make up a service, including

- ServiceManifest.xml
- Service.exe
- Service.dll



SDM Runtime

- **SDM runtime is responsible for tracking SDM models and running instances**
 - **Model tracking**
 - **Model execution**
 - **Model deployment**
- **Components perform operations using the SDM Runtime such as**
 - **Model deployment**
 - **Model execution**
 - **Model tracking**
- **Highly available Instance store (using Yukon's redundant database technology)**
 - **Model execution**
- **Security and account management using Active Directory**

Example: Component Instantiation

```
1 // M.cpp : Defines the entry point for the application.  
2  
3 #include "stdafx.h"  
4 #include "M.h"  
5  
6 #include "MyService.h"  
7  
8 int APIENTRY WinMain(HINSTANCE hInstance,  
9                     HINSTANCE hPrevInstance,  
10                     LPSTR lpszCmdLine,  
11                     int nCmdShow)  
12 {  
13     // TODO: Place code here.  
14  
15     // Create instances of the components.  
16     MyService *fe = new MyServiceFE(hInstance);  
17     MyService *be = new MyServiceBE(hInstance);  
18     MyService *sql = new MyServiceSQL(hInstance);  
19  
20     // Create the catalog.  
21     Catalog *cat = new Catalog(hInstance);  
22  
23     // Create the http server.  
24     HttpServer *http = new HttpServer(hInstance);  
25  
26     // Create the application.  
27     MyCLRApp *app = new MyCLRApp(hInstance);  
28  
29     // Run the application.  
30     app->Run();  
31  
32     return 0;  
33 }
```

myservice.cs is C# code that uses the SOM API

```
componenttype MyService  
{  
    component MyFrontEnd fe  
    component MyBackEnd be  
    port http = fe.http  
    use TDS {  
        fe.catalog  
        be.sql  
    }  
    implementation MyService  
    MyCLRApp  
}
```

myservice.cdm



SOM instances

Operations Logic

Services
Components
Operations

Best Practices
Log Templates

Operations Logic

Operations Logic is the "Business Logic" of Operations

- Operations Logic is CLR code that captures repeatable patterns encoded as reusable best practices

- `public class OpsLogic {`
- `public static void Main() {`
- `// ...`

- OpsLogic is responsible for the overall operation of a service

- `public class OpsLogic {`
- `public static void Main() {`
- `// ...`
- `} // End Main`
- `}`

- OpsLogic is implemented using MS middle-tier technologies

- `public class OpsLogic {`
- `public static void Main() {`
- `// ...`
- `} // End Main`
- `}`


Repeatable Upgrade Patterns → Operations Logic

- Upgrade is an example of the type of Operations Logic we want to build and ship with BIG

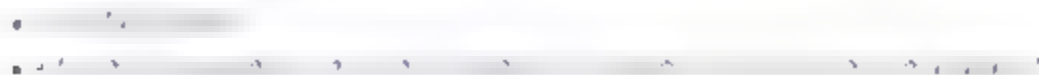
- In-place Upgrade Pattern

- 


- Side-by-side Upgrade Pattern

- 

- Replacement Upgrade Pattern

- 

- Rolling Upgrade is an example of higher-level operations logic that can reuse the codified upgrade patterns

- 

Repeatable Upgrade Patterns → Operations Logic

- Upgrade is an example of the type of Operations Logic we want to build and ship with BIG

- In-place Upgrade Pattern

```
1. Create new version of application
2. Deploy new version to production
3. Verify new version is working
4. Remove old version from production
```

- Side-by-side Upgrade Pattern

```
1. Create new version of application
2. Deploy new version to production
3. Verify new version is working
4. Remove old version from production
```

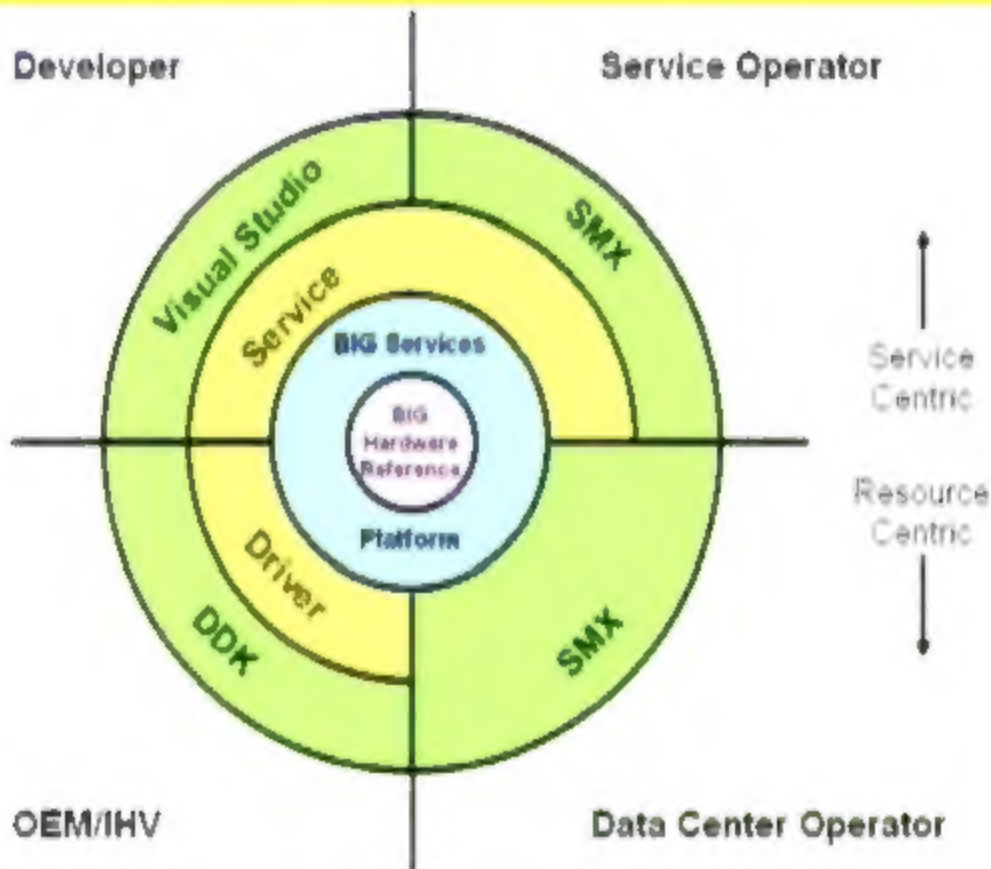
- Replacement Upgrade Pattern

```
1. Create new version of application
2. Deploy new version to production
3. Verify new version is working
4. Remove old version from production
```

- Rolling Upgrade is an example of higher-level operations logic that can reuse the codified upgrade patterns

```
1. Create new version of application
2. Deploy new version to production
3. Verify new version is working
4. Remove old version from production
```


Customer View of BIG



Key Customer Scenarios

- **Enterprise computing utility**

- IT departments as corporate computing utilities to better utilize assets by separating LOB application deployment from hardware procurement
- Example: Goldman Sachs

- **Mega-services**

- Large-scale services that span multiple servers performing parallel or clone functions
- Example: MSN Hotmail

- **Shared-hosting**

- Multiple services hosted on a single server with or without a database
- Example: Schindler or Digeo

- **Diskless network boot**

- Enables a single image to be loaded from the network and run in RAM that is not persisted across reboots to enable centralized image management
- Enables diskless server blades (lower cost/blade, increases reliability and density)
- Example: Morgan Stanley

Schedule

- **Timeline:**

- Current team focused on spec & design of product, building core system and integration with iBIG by end-2002
- Test team and additional dev to roll-off from iBIG at end of 2002

- **M2 – 6/28/02**

- BiG computer built in lab
- Standalone resource manager and drivers running on BiG computer
- Standalone runtime and components (SQL, iIS) on single machine configuration

- **M3 – 8/30/02**

- Integrated network management, resource manager and runtime on BiG computer
- BiG applications using iIS and SQL running on BiG computer

- **M4 – 11/22/2002**

- Integration with iBIG to enable bare metal deployment
- System recovery support

- **RTM – Ship by end of 2003**

Summary

- **BIG proposes a hardware reference platform to lower cost of ownership**
 - BIG Computer simplifies design, deployment and management of distributed, scalable and highly available services
- **BIG is developer-centric, focused on extending the Windows Server Platform for building new services using Visual Studio and reusable building blocks like IIS, SQL, ...**
 - BIG provides technology to build, deploy and operate highly available and scalable services
- **BIG abstracts away the hardware to enable dynamic binding and re-deployment and automated network configuration**
 - BIG delivers infrastructure such as the SDM runtime, resource management, network boot, imaging and file distribution
- **BIG schedule is targeting end of 2003**
 - No dependency on Longhorn